Local Routing in a Tree Metric 1-Spanner

Milutin Brankovic, Joachim Gudmundsson, André van Renssen

March 2020

Preliminaries

Definition

Given a complete, weighted graph $G = (V, E)$, a subgraph H of G is said to be a t-spanner of G if for all $u, v \in V$, there exists a u, v-path in H of total weight at most $t \cdot d_G(u, v)$ where $d_G(u, v)$ is the weight of a shortest u, v-path in G.

A u, v-path in H of total wieght at most $t \cdot d_G(u, v)$ is called a t-spanner path.

Spanners are designed to be sparse, typically with a linear number of edges.

Preliminaries

Definition

Given a complete, weighted graph $G = (V, E)$, a subgraph H of G is said to be a t-spanner of G if for all $u, v \in V$, there exists a u, v-path in H of total weight at most $t \cdot d_G(u, v)$ where $d_G(u, v)$ is the weight of a shortest u, v-path in G.

A u, v-path in H of total wieght at most $t \cdot d_G(u, v)$ is called a t-spanner path.
Spanners are designed to be sparse, typically with a linear number of edges.

Definition

Let T be a weighted tree. The tree metric induced by T, denoted M_T, is the complete graph on the vertices of T where the weight of an edge (u, v) is $d_T(u, v)$.

Clearly, T is a 1-spanner of M_T.

Clearly, T is a 1-spanner of M_T.

Definition

A t-spanner has diameter Λ if every pair of points is connected by a t-spanner path consisting of at most Λ edges.

Clearly, T is a 1-spanner of M_T.

Definition

A t-spanner has diameter Λ if every pair of points is connected by a t-spanner path consisting of at most Λ edges.

Solomon and Elkin address the problem of designing a 1-spanner for tree metrics which simultaneously achieves low diameter, low weight and low degree.

Definition

Given a wieghted graph G, the lightness of G is the ratio of its weight to the weight of a minimum spanning tree.

Definition

Given a wieghted graph G, the lightness of G is the ratio of its weight to the weight of a minimum spanning tree.

Theorem (Solomon and Elkin, 2014)

For any weighted tree T of maximum degree Δ, there exists a 1-spanner for M_T which has

- *diameter* $O(\log n)$
- *lightness* $O(\log n)$
- *maximum degree* $\Delta + O(1)$

Definition

Given a wieghted graph G, the lightness of G is the ratio of its weight to the weight of a minimum spanning tree.

Theorem (Solomon and Elkin, 2014)

For any weighted tree T of maximum degree Δ, there exists a 1-spanner for M_T which has

- *diameter* $O(\log n)$
- *lightness* $O(\log n)$
- *maximum degree* $\Delta + O(1)$

We show that a slightly simplified version of this spanner supports a local routing algorithm which takes advantage of the low diameter.

Local Routing

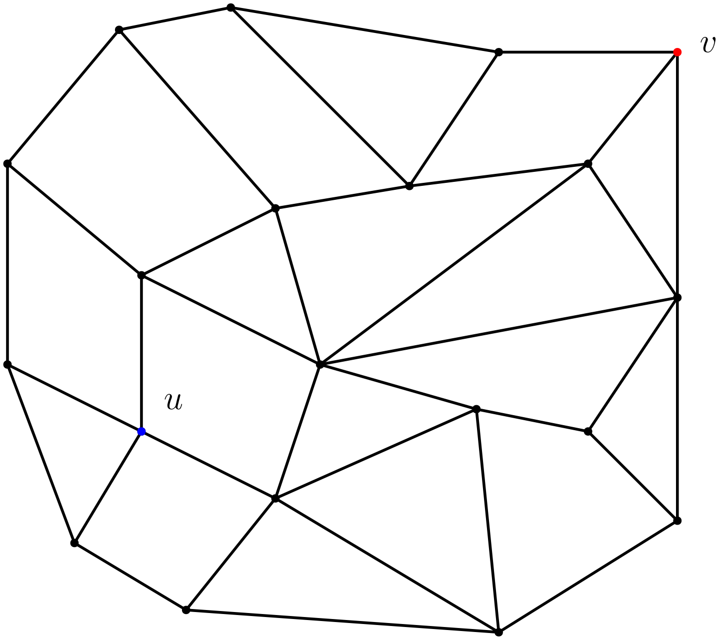

Local Routing

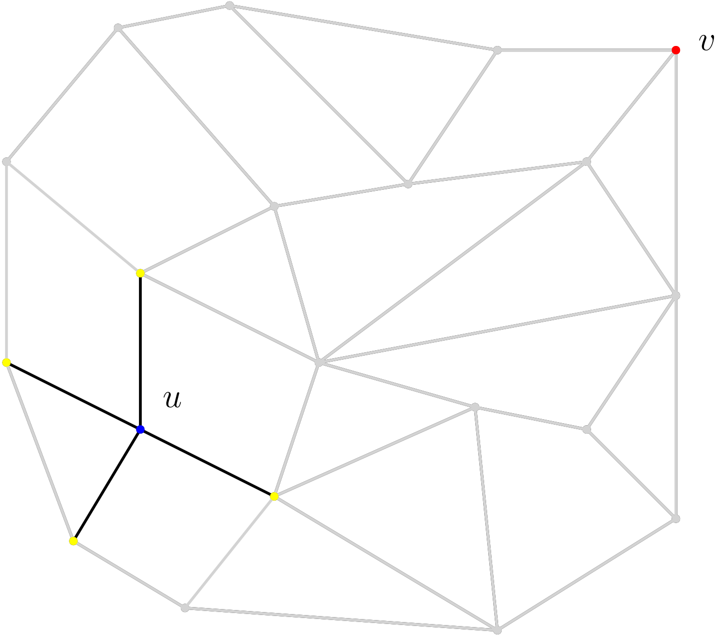

Local Routing

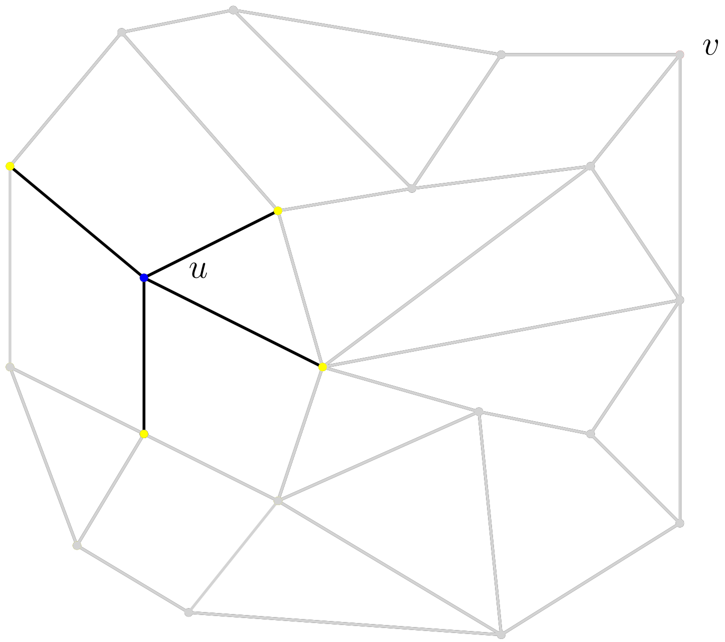

Local Routing

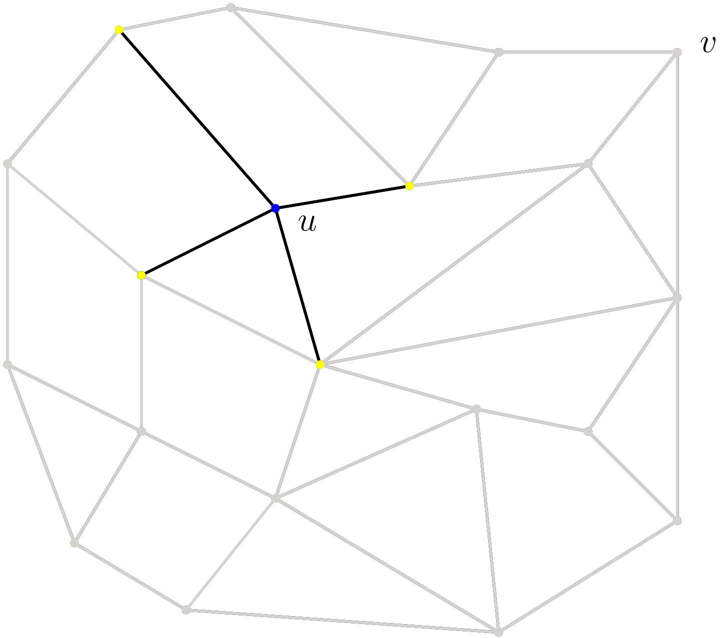

Local Routing

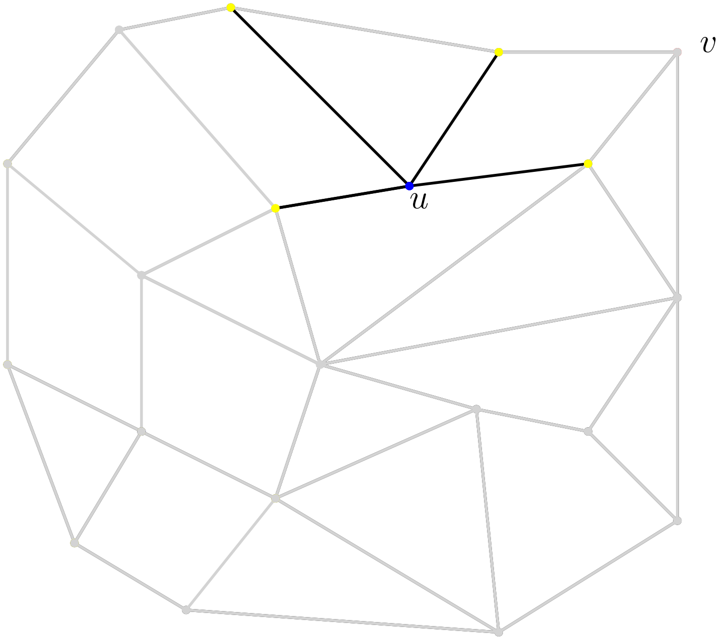

Local Routing

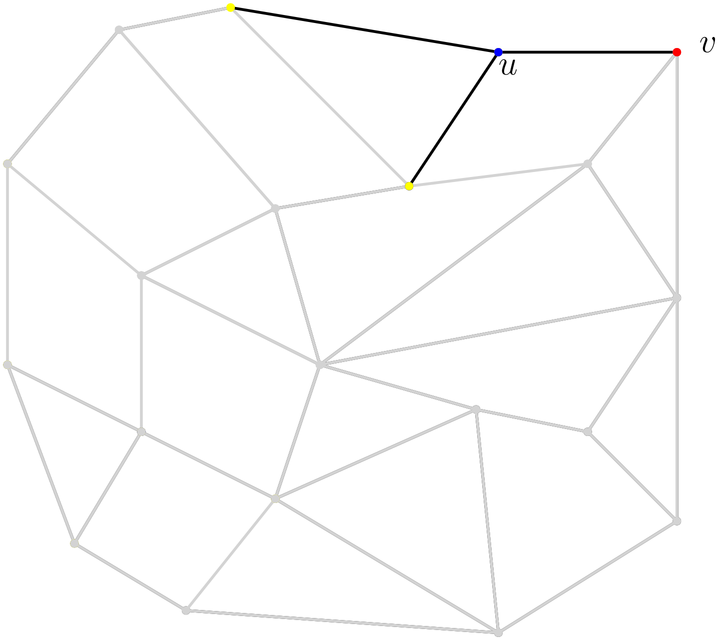

Local Routing

Definition

A local routing algorithm has diameter Λ if it is guaranteed to terminate after traversing Λ edges.

Local Routing

Definition

A local routing algorithm has diameter Λ if it is guaranteed to terminate after traversing Λ edges.

Definition

Given a local routing algorithm on a weighted graph $G = (V, E)$, let $d_{\text{route}}(u, v)$ denote the weight of the path traversed when routing from u to v. The routing ratio of the routing algorithm is defined to be

$$\max_{u,v \in V} \frac{d_{\text{route}}(u, v)}{d_G(u, v)}.$$

Labelling Scheme for Trees

For a vertex v in T, let $rank(v)$ denote the rank v in a post-order traversal of T.

Labelling Scheme for Trees

For a vertex v in T, let $rank(v)$ denote the rank v in a post-order traversal of T.

Define $\min(v) := \min\{rank(w) : w \text{ is a descendant of } v\}$.

Labelling Scheme for Trees

For a vertex v in T, let $rank(v)$ denote the rank v in a post-order traversal of T.

Define $\min(v) := \min\{rank(w) : w \text{ is a descendant of } v\}$.

The label of v is defined to be

$$\ell(v) := [min(v), rank(v)].$$

Labelling Scheme for Trees

For a vertex v in T, let $rank(v)$ denote the rank v in a post-order traversal of T.

Define $\min(v) := \min\{rank(w) : w \text{ is a descendant of } v\}$.

The label of v is defined to be

$$\ell(v) := [min(v), rank(v)].$$

Observe that for two vertices u and v, u is a descendant of v if and only if $rank(u) \in [\min(v), rank(v)]$.

Local Routing in Trees

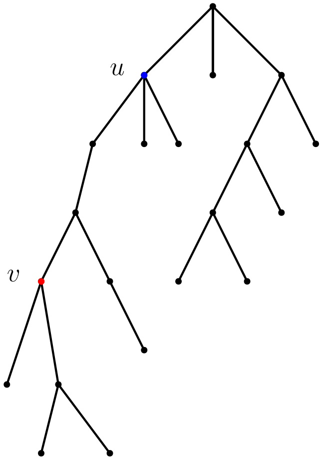

Case 1: u is an ancestor of v.

Route to the child of u which is an ancestor of v.

Local Routing in Trees

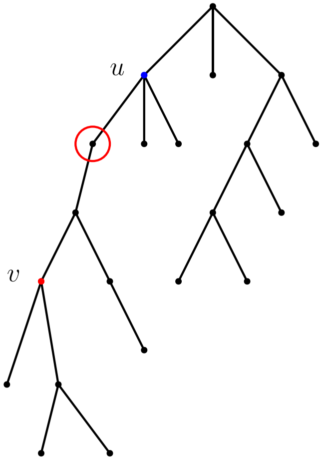

Case 1: u is an ancestor of v.

Route to the child of u which is an ancestor of v.

Local Routing in Trees

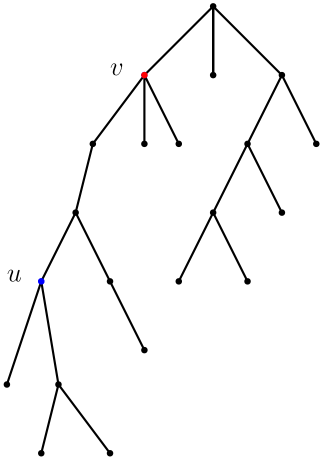

Case 2: u is not an ancestor of v.

Route to the parent of u.

Local Routing in Trees

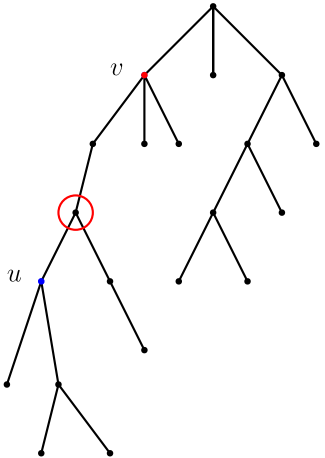

Case 2: u is not an ancestor of v.

Route to the parent of u.

Local Routing in the 1-Spanner

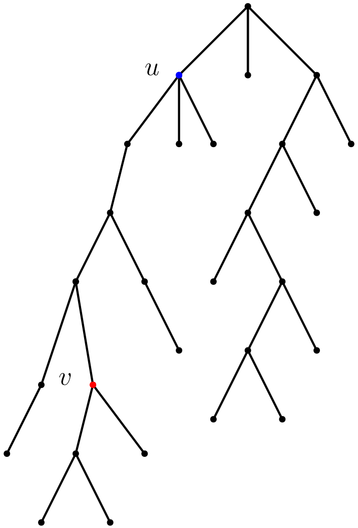

Case 1: u is an ancestor of v.

Local Routing in the 1-Spanner

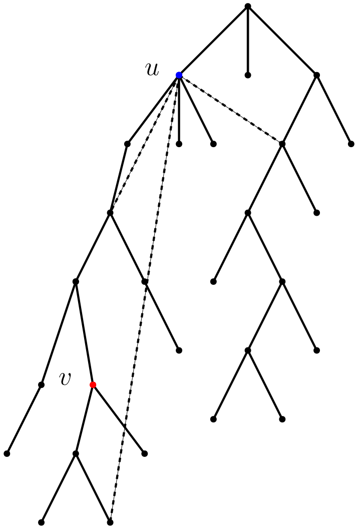

Case 1: u is an ancestor of v.

Route to the deepest neighbour of u which is an ancestor of v.

Local Routing in the 1-Spanner

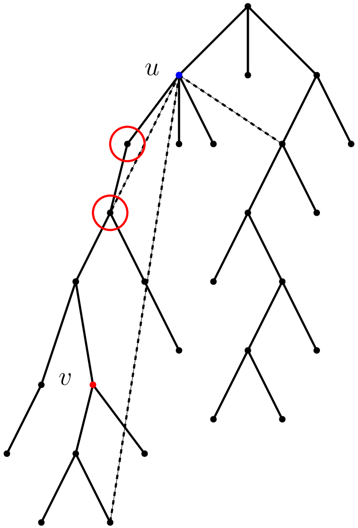

Case 1: u is an ancestor of v.

Route to the deepest neighbour of u which is an ancestor of v.

Local Routing in the 1-Spanner

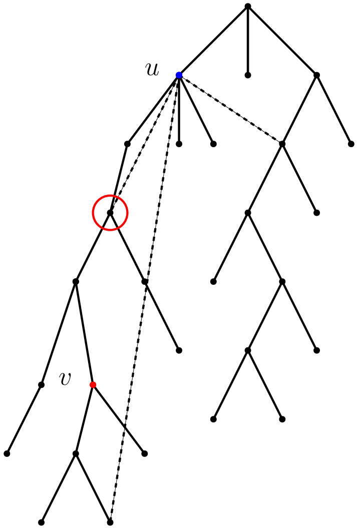

Case 1: u is an ancestor of v.

Route to the deepest neighbour of u which is an ancestor of v.

Local Routing in the 1-Spanner

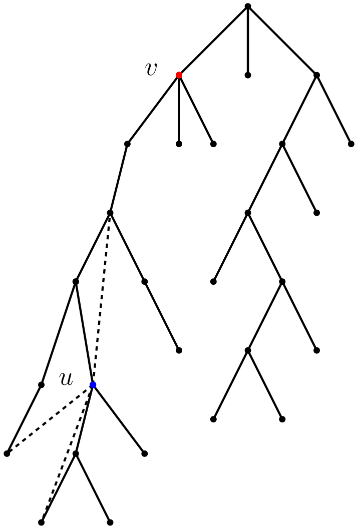

Case 2: u is an descendant of v.

Local Routing in the 1-Spanner

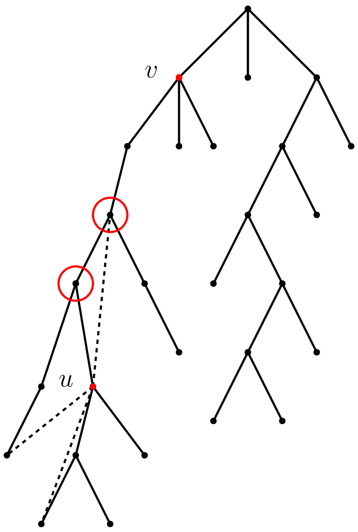

Case 2: *u* is a descendant of *v*.

Route to the highest neighbour of *u* which is a descendant of *v*.

Local Routing in the 1-Spanner

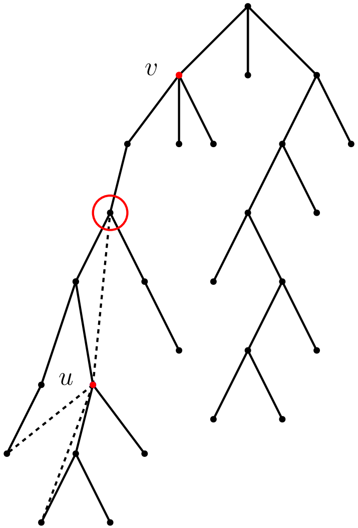

Case 2: *u* is a descendant of *v*.

Route to the highest neighbour of *u* which is a descendant of *v*.

Local Routing in the 1-Spanner

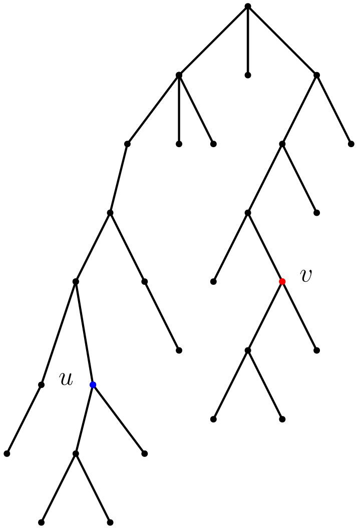

Case 3: u is not a descendant or an ancestor of v.

Local Routing in the 1-Spanner

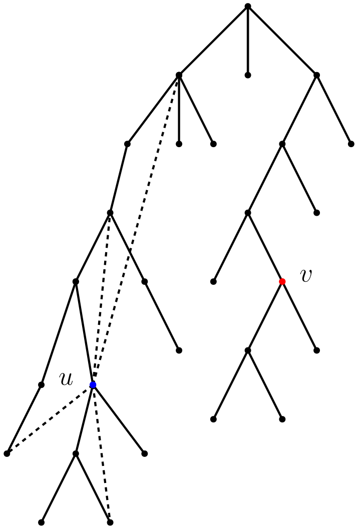

Case 3: u is not a descendant or an ancestor of v.

Case 3a): No neighbour of u is an ancestor of v.

Local Routing in the 1-Spanner

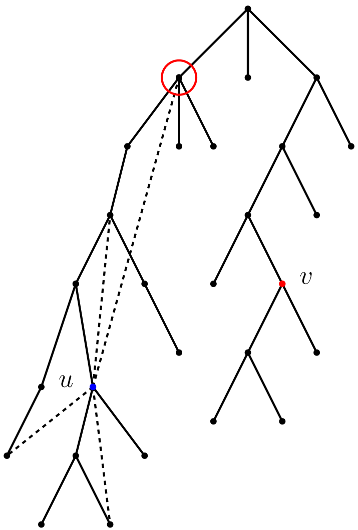

Case 3: u is not a descendant or an ancestor of v.

Case 3a): No neighbour of u is an ancestor of v.

Route to the highest neighbour of u.

Local Routing in the 1-Spanner

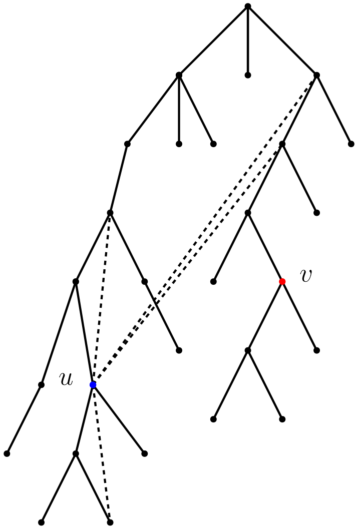

Case 3: u is not a descendant or an ancestor of v.

Case 3b): u has neighbours which are ancestors of v.

Local Routing in the 1-Spanner

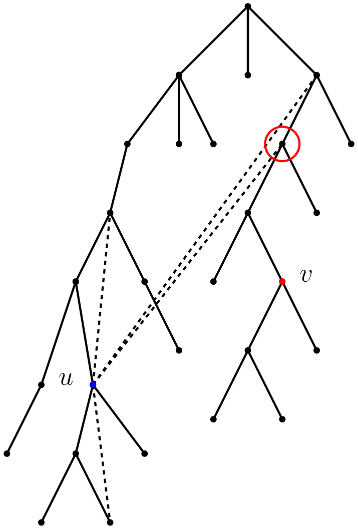

Case 3: u is not a descendant or an ancestor of v.

Case 3b): u has neighbours which are ancestors of v.

Route to the deepest neighbour of u which is an ancestor of v.

Conclusion

We can show that when applied to a simplified version of the 1-spanner, the simple routing algorithm we have described is guaranteed to terminate after $O(\log n)$ steps.

Conclusion

We can show that when applied to a simplified version of the 1-spanner, the simple routing algorithm we have described is guaranteed to terminate after $O(\log n)$ steps.

Theorem

A simplified version of the 1-spanner construction of Solomon and Elkin supports a $O(\log n)$*-diameter local routing algorithm with routing ratio 1.*